

An Interview with
WILLIAM CROWTHER
OH 184

Conducted by Judy O'Neill
on
12 March 1990
Cambridge, MA

Charles Babbage Institute
Center for the History of Information Processing
University of Minnesota, Minneapolis

Copyright, Charles Babbage Institute

William Crowther Interview
12 March 1990

Tape Index

Tape 1/Side 1.....	1
Side 2.....	17

WILLIAM CROWTHER INTERVIEW

DATE: March 12, 1990

INTERVIEWER: Judy O'Neill

LOCATION: Cambridge, Massachusetts

O'NEILL: Let's start with a brief sketch of your career, what your educational background is, and your work experience before starting on the ARPANET.

CROWTHER: Sure. I went through MIT, got a B.S. in 1958. Then I went to work for MIT at Lincoln Lab for about ten years. I came to BBN for about eight. Went off to Xerox PARC for another seven or eight, and then came back to BBN again. I've been here for five or six, I think it is now. You'll find that my grasp of times and history and such is pretty fuzzy. All these are plus or minus two or three years.

O'NEILL: Well, hopefully we'll be able to step through the parts we need. What was your B.S. degree in?

CROWTHER: Physics. In those days they didn't have a computer science department. In fact, the computer wasn't even in the electrical engineering department; it was in the physics department, I think, because they were the only people who were going to try and keep it running.

O'NEILL: Which computer was that, do you remember?

CROWTHER: Well, it was an old IBM something or other, 704... it's got to be something earlier than 704. One of those things.

O'NEILL: Can you describe some of the kinds of systems that you worked on while you were working at Lincoln Labs?

CROWTHER: Yes. Let's see, they tended to be real-time control systems. There was a thing that pointed one of these large infrared antennas that MIT had at that time. There was another thing, a truck that was doing communications, bouncing signals off the moon or passing satellites, that kind of thing. Lincoln Lab tended to do state-of-the-art communications kinds of things. They were trying to make a mobile truck to replace the big fixed ground stations that the government had at that time. The mobile truck worked great, and it didn't replace anything. [Laugh]

O'NEILL: Were you interested in computers right away?

CROWTHER: Yes. My thesis was something to do with computers. I have actually sort of forgotten. It had to do with primal dual method of solving sets of simultaneous inequalities. Sort of related to the simplex thing. It was a mess. [Laugh] But that was a B.S. thesis; those aren't very fancy anyway.

O'NEILL: Did you work on something called the Lincoln Experimental Terminal System? Can you explain what that was?

CROWTHER: Yes. That was that truck. It had a small computer in it, and it had, actually, liquid nitrogen cooled electronics at the heart of the antenna, which nobody did in those days. My part in it was to make the computer do its tricks.

O'NEILL: How was the Lab structured at that time? What group were you working in at Lincoln?

CROWTHER: That's a good question. I don't really remember all that administrative stuff.

O'NEILL: How about the people you were working with?

CROWTHER: Well, Frank Heart was prominent in most of the things that I did. I liked to work for Frank. So he was one of the key figures in my life. He was just a little bit older than I was, and he tended at that point to be running projects, and I tended to be writing computer programs, loving the fact that people paid me for playing.

O'NEILL: Was Frank a technical person as well as administering the projects?

CROWTHER: Well, Frank had technical control. When he ran a project, he wouldn't let go of anything until he completely understood every little piece of it. So he, in fact - through the ARPANET, too - he knew everything that was going on in the technical part even though he didn't actually implement anything. This was, I thought, a very good thing because it meant everybody had to explain everything to Frank, and by the time he understood it, everybody else understood it, too.

O'NEILL: Were you aware of or interested in the work that Larry Roberts and Tom Marill were doing, connecting the SDC Q32 and the TX2 at Lincoln Labs?

CROWTHER: Well, I knew they were doing it. When the ARPANET started, I knew that Larry was doing that. I guess I sort of knew he was doing it even at Lincoln. I didn't pay much attention to that, actually. I tend not to pay attention to anything except what I'm doing. In those days, the thing I cared most about was rock climbing, so... [Laugh]

O'NEILL: And that fitted in with your antenna research?

CROWTHER: Well, antenna research paid for the rock climbing. [Laugh] I knew he was there. There were lots of interesting things happening at Lincoln. And you sort of kept on top of what was going on, just because it was fun to do. But I didn't know too much about what Larry was doing.

O'NEILL: But you did know him?

CROWTHER: Yes. We went skiing once, I think. Yes, I knew him.

O'NEILL: Can you tell me approximately when you first came to BBN?

CROWTHER: I tell you, I'm bad at dates, but it had to be just before the ARPA project started, just before the proposal. They basically hired me when they thought they would have work to do on the ARPANET.

O'NEILL: That would have been about 1968 then.

CROWTHER: That sounds right, maybe 1967.

O'NEILL: Why did you come?

CROWTHER: Oh, because I thought it would be great to work for Frank. And there is this funny thing that happens. Organizations tend to get old. More exciting things were happening at BBN, or at least they were happening in a more exciting way, than they were at Lincoln Lab. So I kind of enjoyed that switch.

O'NEILL: Do you remember any of the other people who had already moved over to BBN that you knew from Lincoln?

CROWTHER: Well, Severo had. I forget exactly what the order was. I guess Dave came first; I'm not sure.

O'NEILL: Yes. I guess I wasn't sure how much you had worked with him at Lincoln.

CROWTHER: I worked with Dave a lot. I sort of helped Dave write his first program at Lincoln. That was kind of fun.

O'NEILL: When you got to BBN, did you work on the proposal for the response to the RFQ?

CROWTHER: Yes.

O'NEILL: At that time, did you start investigating networking at all? It does not seem like you were really into computer networking per se before that. You were working on the real-time systems.

CROWTHER: This looked like just another real-time system - one of these things that you have to get your head around. After you see how all the pieces fit together, it wasn't much different from pointing an antenna or doing other kinds of things.

O'NEILL: So is it fair to say that your view of networking, or of your project anyway, at the time was that it was just another real-time system, similar to what you had done before at Lincoln?

CROWTHER: A good complicated one. I like complicated ones.

O'NEILL: I think the timing is going to be off here, but I'll go ahead and ask anyway. Larry Roberts at ARPA started having meetings, in preparation for the RFQ. Would you have had any involvement with that?

CROWTHER: No, I didn't. I think Bob Kahn was in on that. But I was not.

O'NEILL: Do you remember having any knowledge of other people working on these kinds of networks, like the work of Paul Baran at Rand, or of Donald Davies at NPL?

CROWTHER: No. I guess the ones I knew about were, well, I knew that Larry had done some little thing, and also that someone in England - that's all I remember about who - had implemented something. That one was interesting because we were projecting doing things quite a bit faster than they had, at least an order of magnitude faster than they were doing, which was disturbing to a number of people.

O'NEILL: When you say disturbing, in what way? Did people come out and say, "You can't do this?"

CROWTHER: Yes, basically. Our response to that was, "Yes we can; we've coded the kernel of the thing, and we know how fast it's going to run, and it can indeed process ten times as many packets per second as this other system."

O'NEILL: Who were the people saying that you couldn't do it? I don't mean necessarily individuals but the types of people. Were they people from the telephone company, or people doing academic research?

CROWTHER: Well, I don't remember that too well. I think it was the academic sorts and the people at DARPA. We had to justify that we could actually make this performance to the people at ARPA. I forget and call it DARPA every once in a while.

O'NEILL: We go back and forth all the time.

CROWTHER: There was considerable skepticism at first. I don't exactly remember where it was coming from.

O'NEILL: But there was voiced skepticism that it was not going to work?

CROWTHER: Well, that it was not going to work at the speeds that we were saying. In fact, it wouldn't have had to work at the speeds to be acceptable to them. The original RFP specified a thing that was ten times slower than what we actually did.

O'NEILL: When were you convinced that it was going to work ten times faster? Were you convinced from the beginning?

CROWTHER: What actually happened was Dave and I sat down, worked out the algorithms, figured out that it was only going to take a hundred and fifty lines of code to process a packet through one of these switches. We actually sat down and wrote the hundred and fifty lines of code, and counted them, and then we knew.
[Laugh]

O'NEILL: Did you actually have a machine to run that on?

CROWTHER: No. We had no machine. It certainly wouldn't have worked, because there's a whole lot of stuff outside the inner loop that you have to have to maintain the state. But the actual thing took in a packet, figured out what to do with it, and pushed it back out the line. It was very short and quite practical, too. We knew exactly what it was going to do.

O'NEILL: And that was based on your experience with other real-time systems?

CROWTHER: Yes. That is how you figure out if real-time systems are going work. You write the kernel, usually there is some very small part that is the only thing that matters; and once you have that one figured out, you know what the timing is going to be.

O'NEILL: Did you have any trouble when you actually had to put that onto a specific computer?

CROWTHER: No. Except for the hardware troubles that they had. Surely Severo has told you about those.

O'NEILL: Yes, he mentioned a few.

CROWTHER: Did he mention that the hardware was finally working just a few days before we had to deliver? And we had constructed simulators so that we would have a chance to run our software and at least debug it in the simulated state?

O'NEILL: No. I didn't realize that.

CROWTHER: Yes. The software actually... There was almost no time in which to debug it on the real machine, and we made it work anyway.

O'NEILL: Had you planned all along to have simulators available so that you could test it out?

CROWTHER: No.

O'NEILL: So that was a stop-gap measure when the hardware wasn't coming along as scheduled?

CROWTHER: [Yes.] Well, it was practical certainly. The whole program was pretty small. As I recall, in those days it was 4K, but that must be words, so it's 8K bytes the whole thing fit in, and that included buffers, too, so it was a pretty small program. It was practical to hand de-bug it and get it mostly right and have the simulators check it out, too.

O'NEILL: Was using simulators a standard technique at the time?

CROWTHER: Not Ben. The others had. I had worked with all of them before on systems that seemed to be very much like this one. It was a really good group - made sensible things happen.

O'NEILL: What made this project interesting, if it was like what you had done before?

CROWTHER: Oh, it was just another one - another fun puzzle. I was willing to do an almost unlimited number of those. [Laugh]

O'NEILL: It sounds like a pretty small group, so this may be an obvious question, but did you work really closely with the people working on the hardware? Were you aware of all the hardware problems, and the progress on the hardware side?

CROWTHER: Even more than that, they would give a complete description of their design and I would sit there and say "No, no this isn't the best way." And vice versa, we would completely describe the software and they would sit in there and say, "No, why are you doing it this way?" Yes, everybody knew everything.

O'NEILL: So there was a lot of interaction on a daily basis?

CROWTHER: Yes. Severo sat in the office next to mine.

O'NEILL: What was your role in the actual implementation? I know Dave Walden talked about going out to sites and installing the tapes when the first sites were coming on line.

CROWTHER: I did some of that, too.

O'NEILL: Can you describe what that was like? Was it fairly straightforward implementation? Were there a lot of problems? Were the sites easy to work with?

CROWTHER: I only did it, as I recall, when we were bringing up the first four node network. And it was kind of fun. And you get to see new people, and you get to try to make the thing work. Mostly it did work, so that wasn't so bad. I don't have any striking memories from that time. I think we spent most of our time trying to figure out how to help the people who were trying to get their host programs communicating, giving them some type of clue as to where things were going wrong.

O'NEILL: Do you remember which sites you actually went to?

CROWTHER: I don't remember. Utah sticks in my head, but that is just because I have visited since. When the original thing came out, we were trying to have one of us at each site. I just don't remember which site I went to.

O'NEILL: Once the network got installed, what was your role then? Did it change significantly, once you got the four nodes up and running? Did you see a shift in what you were doing?

CROWTHER: I am trying to remember. We were either adding new features, fooling around with the routing and such. Initially we were all involved in the day-to-day thing when you added new nodes. But that quickly got turned over to some other people. Tony Michel and Kotzky were involved at that time. I don't know exactly when we started working on the terminal concentrator. But that is what sticks in my head as the next thing we did, to make that piece.

O'NEILL: I was coming up with the terminal IMP - that is what you are referring to?

processor implementation of the IMP or a terminal IMP would be a pretty good idea. So that is how the Pluribus project got started.

O'NEILL: So it was started in response to a requirement from ARPA, but not a recommendation of how to meet that requirement?

CROWTHER: They didn't care how it was done. Both Severo and I thought it would be fun, and so Frank thought it would be fun, so he was pushing it. It had a lot of advantages. It had some very interesting reliability aspects to it. I don't know whether you talked to Severo about that thing, but there was a piece of code in there, most of which I wrote, where you could pull out any card or pull out any wire, or short out any component (except we didn't like to do that very often) and the thing would keep running. Mostly the program would keep running. So it would keep on behaving as a terminal concentrator.

O'NEILL: So it was fault-tolerant?

CROWTHER: It was fault-tolerant in a funny way. That is, you often think of fault-tolerance as the way the banking systems do, where you're not willing to accept an error. But in a communications system what you really care about is if the thing is still up and running. You don't care whether it fails to deliver a message because all sorts of things could cause it to fail to deliver a message. So it really was designed to stay up and running, rather than not making any errors. For example, there was a background task that went around and looked at all of the buffers in the machine to see that they were on some queue. If they weren't on some queue, then it would pick them up and put them back on the free list. No matter what happened, hardware failure, software failure, you couldn't really run out of buffers ever. [Laugh]

O'NEILL: Because they would always be picked up, no matter how they got disassociated.

CROWTHER: Right. There was a lot of stuff like that.

O'NEILL: When problems came up, or new parts of the network, things like the various approaches to routing, did you actually do experiments? Did you just think it through? Did you use theory? How did you go about coming up with them?

CROWTHER: All of those, all of those. If given a complex system and an algorithm, like a routing algorithm, I tend to be pretty good at visualizing the thing and seeing what will happen and what some of the bad cases are. So there were a lot of mental things like that. When you came up with one that looked pretty good, then you'd try it and see whether or not it worked.

O'NEILL: Did you have a test network that you could play with? How did you go about trying these?

CROWTHER: Well, not originally. There were always a few machines in the back room that were being built, and before they got completely checked out and shipped, we could throw together little three and four node networks to make it go. We tended, I tend anyway, to favor the simple algorithms. They may not work wonderfully, but they're probably not going to break terribly, either. I am sure you have got people who have told you about the terrible things that happen when some node says, "I am the route to everywhere."

O'NEILL: I've seen that mentioned in some articles.

CROWTHER: No one had thought of that one before. But all it took was one hardware failure, and that happened, and that brought the net down.

O'NEILL: So that would be an example of a surprise.

CROWTHER: Yes, that was quite a surprise. And that led to this notion of distributing the system through the network in such a way that you could isolate yourself. You could force a spread from machine to machine even if the machine at the other end of the line wasn't listening to you properly. You could force him to reload and ignore his neighbors who weren't behaving right.

O'NEILL: Did that just become obvious as you started working on this, that that was a good approach to take?

CROWTHER: Well, the thing that made it... I think, it went over a cliff when one of the crucial nodes, cross country link, was in a military base, and it broke on a Friday evening. And we couldn't get into the base until the next Monday. That was the straw that made us go to these things where we had to be able to do it without any access to the site.

O'NEILL: Were there other people writing about these subjects? Did you actually go out and do research on how to do this or was it just a matter of "We've got a problem here that we need to fix."

CROWTHER: Mostly it was, "We have a problem, we need to fix it, and we're certainly willing to listen to any clever ideas." But there were all sorts of crazy ideas about, and most of them didn't make any sense. There was this 'hot potato' routing which somebody was advocating, which was just crazy. There were whole lots of algorithms that just didn't make any sense.

TAPE 1/SIDE 2

CROWTHER: So mostly what we did was we would steal ideas from anywhere, but most of the time we had to roll our own.

O'NEILL: What was your interaction was with the rest of the community, the people at the host sites, for instance. Did you interact with them a lot? Did you present your ideas to them?

CROWTHER: Well, that depends a lot about when in the project you are talking.

O'NEILL: Let's start early.

CROWTHER: Early on Frank had made a decision that I think was a very wise decision to make a clean boundary between the host responsibilities and the network responsibilities. So for a while we were focusing only on our problems - delivery of the messages. Anything that was host to host was someone else's responsibility, and we weren't going to be concerned with that. When our stuff started to work, and when we started to do the terminal things, we actually were a host, then we got more involved in these protocols. We would go to some of the protocol meetings and listen to the people and talk to the people. So I knew people like Vint and some of the others doing this kind of thing. Mostly I thought it was dull. [Laugh]

O'NEILL: Having the meetings or talking about protocol?

CROWTHER: Both. I tend to think that these network things are actually pretty simple. I know that there's a whole industry out there that has developed off of networking and layers and layers of protocols and all that kind of stuff. But, in fact, if you just sit down from scratch and try to build one of these things there is a natural way to do it, and it does have layers because of course any good programming will have layers. Mostly the natural things work. There seemed to be two things going on in the protocol meetings that were not very productive. One was that some people would have done it one way and some other people have done it another way, and either way would work. And there would be great fights trying to get one side or the other to give in. Then there would be a third thing, which was someone would be trying to come up with something that was

novel, probably because they wanted to get a paper out of it. So you would have to cope with that kind of thing.

O'NEILL: Can you give me an approximation of how large these meetings were when you were going to them? Was it ten people sitting around in a room? Was it 200?

CROWTHER: Oh, I have no idea. When they got big, I quit. They tended to be smaller numbers - ten, fifteen.

O'NEILL: Let's just get back a little bit to your career. You left BBN and went out to Xerox PARC in about 1976. Did you continue working in networking areas or real-time systems?

CROWTHER: For about a year. Xerox at that time thought it was going to cover the world with workstations to make the paperless office. And they needed communications for that, so they asked me to come out and I built the communications sub-system for their STAR, it was called in those days the Star System. After a little while it became clear that nothing was going to happen. The politics were too horrible, the technical decisions were being made by politicians, all that kind of stuff. And it wasn't going to turn into a real product. So I decided I was wasn't going to stay there. What I did do was move across the street to the research part of PARC, and basically got out of the communications stuff and started doing other things.

O'NEILL: By this point in your career did you see networking as separate from other real-time systems? It sounds like you were brought to PARC for your communications expertise. Is that how you viewed yourself as well?

CROWTHER: No. I tended to think of myself as someone who could write almost any kind of program that was tricky.

O'NEILL: When you were being encouraged to disseminate more information, what was the general reaction to that here at BBN - you, personally?

CROWTHER: Well, writing papers is less fun than implementing things. But it is also important, so I eventually did it. The ARPANET was fairly well written up in the long run. It came a little bit slowly.

O'NEILL: Was there also encouragement for things like going to conferences and giving presentations?

CROWTHER: Yes.

O'NEILL: Did you have any exposure to the military during this time?

CROWTHER: I don't think the military was involved. Well, let's see. In the early stages it was all universities. Eventually some military nodes got attached. But there wasn't much interaction there. It tended to be that site A would want to send the mag tape to site B every day - they had been putting it in the airplane every day for the last year. And Larry said, "This is silly. Let's hook you onto this network we have, and use some of our excess capacity." So they hooked on. They were not terribly excited about it, but it did save them putting the thing on an airplane.

O'NEILL: Was that the justification for doing the magnetic tape option on the IMP?

CROWTHER: I guess. Also it seemed like a good idea.

O'NEILL: So you don't remember if there was a particular military request for it?

CROWTHER: No I don't. I don't remember that. It sounds familiar now that you mention it, but if you had asked me cold, I would not have come up with it.

O'NEILL: We've gone through quite a bit. The story of the ARPANET has been told a few times already. I am assuming you're familiar with some of the historical accounts.

CROWTHER: I read some of them.

O'NEILL: Is there anything that you disagree with, or that you would state differently from what is generally known about the development of the ARPANET?

CROWTHER: I don't think so. Every account you read stresses different players in different ways. Well, that's fair. We did our part, and our part was fun. Other people did their parts, and they get to stress that, too.

O'NEILL: Do you have any other general comments on your involvement with the ARPANET? Anything that you would like to add?

CROWTHER: As I've said lots of times, I thought it was fun. That, for me, is an important criterion in what I'm doing. So... I liked it. It was nice that it got used, that it became a real thing. It makes my resume read better. [Laugh] But most importantly, I had a good time building it.

O'NEILL: Did you get back into networking when you came back to BBN after being at PARC?

CROWTHER: Well, BBN is a funny place. And I am a principal scientist at BBN right now. That means that I get involved in everything. But I've stayed out of the networking mostly. There are some people working on a high speed gigabit switch. I talk to them every once in a while, because that is sort of a fun thing. When I